Langevin Diffusion Transitional Markov Chain Monte Carlo with an Application to Pharmacodynamics

Georgios Arampatzis, Daniel Wälchli, Panagiotis Angelikopoulos*, Stephen Wu†
and Petros Koumoutsakos

Professorship for Computational Science, ETH-Zurich, CH-8092, Switzerland

October 19, 2016

arXiv:1610.05660v1 [stat.CO] 18 Oct 2016

Abstract

We propose an algorithm for the sampling of the posterior distribution in Bayesian inference problems. The algorithm builds on the Transitional Markov Chain Monte Carlo (TMCMC) method with independent MC chains following the manifold Metropolis Adjusted Langevin transition kernels (mMALA). The algorithm combines the advantages of population based sampling algorithms – global exploration of the parameter space, parallelizability, reliable estimator of evidence – with the advantages of Langevin diffusion based Metropolis-Hastings MCMC – efficient exploration of local space. Here, we use a bias free version of the TMCMC algorithm, namely the BASIS algorithm. We show that the proposed sampling scheme outperforms the existing algorithm in a test case of multidimensional Gaussian mixtures and the posterior distribution of a Bayesian inference problem in an 8 dimensional pharmacodynamics model.

1 Introduction

Bayesian inference has been widely used for parametrizing mathematical models of physical process, combining observed data, intuition and expert knowledge [14] . Using the Bayes theorem, $p(\vartheta|\mathcal{D}, M) \propto p(\mathcal{D}|\vartheta, M)p(\vartheta|M)$, we are able to quantify our degree of belief on the values of parameters ϑ of a physical process described by model M, conditioned on the observations $\mathcal{D}$. The benefit of the Bayesian framework is that it provides the full posterior probability of the parameters and thus the uncertainty in the predictions of the model can be robustly quantified.

The sampling of $p(\vartheta|\mathcal{D})$ is done using numerical techniques and the quality of the samples relies on the ability of these algorithms to explore the, in most situations high dimensional, parameter space. A well studied and broadly used class of algorithms for the sampling of a general, potentially un-normalized, probability distribution is that of the Metropolis-Hastings (MH) algorithms [20]. These algorithms are a subset of the general framework of Markov Chain Monte Carlo (MCMC) algorithms. Given the target distribution, these algorithms provide a Markov sequence, $\vartheta_1, \ldots, \vartheta_N$, of samples that converges to the target distribution as a stationary distribution. Conditional on ϑ_N, a new sample $\vartheta^\star$ is proposed according to a probability distribution $q(\cdot|\vartheta_N)$ and the sample is accepted or rejected with probability $\alpha(\vartheta^\star, \vartheta_N)$ given by (9). The proposal probability can be freely chosen and the convergence properties of the algorithm heavily hinge upon this choice. A common choice is the normal distribution centered at ϑ_N and with a constant covariance matrix. A poorly chosen covariance matrix can lead to poor quality samples; if the covariance is large then most of the proposal samples will be rejected, and if the covariance is small, most of the samples will be accepted and the probability space will not be sampled. For a detailed discussion on this issue we refer to [23].

An improved sampling algorithm is offered by the Metropolis Adjusted Langevin Algorithm (MALA) proposed in [21, 22]. In this algorithm the proposal distribution is derived from the discretization of a Langevin diffusion with a drift term related to the gradient of the target distribution. The remaining issue of the optimal selection of the covariance matrix is an open problem.

*currently at D. E. Shaw Research L.L.C, 10036 New York, NY, USA
†currently at Institute of Statistical Mathematics, 190-8562 Tokyo, Japan

The authors in [10] make a connection of the MALA algorithm with a discretization of a Langevin diffusion on a general manifold. Then, they propose a covariance matrix based on the Hessian or the Fisher information of the target distribution. The use of a covariance related to the Hessian was also proposed in [6] through the point of view of optimization using the Newton method. Another class of algorithms with the ability to provide proposals with high acceptance probability are the Hybrid or Hamiltonian Monte Carlo (HMC) [9].

Regardless of the popularity of the MCMC algorithms, they usually lack the ability to explore multi-modal or highly peaked parameter space. These algorithms can be trapped indefinitely long time in local modes. An efficient class of algorithms to tackle such issues is the Parallel Tempering or Population MCMC. Here, we consider a special case of these algorithms for the sampling of posterior distributions, namely the Transitional Markov Chain Monte Carlo (TMCMC) algorithm proposed in [8]. An improved, unbiased version of this algorithm, called BASIS, was proposed in [25]. For a review and current trends in Bayesian computations we refer to [11].

In this paper we propose and implement a combination of the TMCMC algorithm with Langevin diffusion transition kernels by following the manifold approach presented in [10]. We use both, Hessian and Fisher information to construct the covariance matrix of the proposal distribution and discuss the advantages and disadvantages of each choice. Finally, we show that the proposed algorithm is able to sample the posterior distribution in a case where the TMCMC fails.

The paper is organized as follows: in Section 2 we give basic background information on Bayesian inference, TMCMC and manifold MCMC, in Section 3 we present the manifold TMCMC algorithm and discuss the implementation details, in Section 4 the application of the proposed algorithm in ordinary differential equations (ODE) is discussed and in Section 5 we show results of our algorithm in a 10 dimensional Gausssian mixture probability distribution and the posterior of an 8 parameter pharmacodynamics ODE system.

2 Background

In this section we give a brief description of the Bayesian inference problem, the Transitional Markov Chain Monte Carlo (TMCMC) algorithm [8], and the manifold Metropolis Adjusted Langevin Algorithm (mMALA) proposed in [10].

2.1 Bayesian Inference

Given a data set $\mathcal{D} = \{d_i \,|\, i = 1, \ldots, N_{\mathcal{D}}\}$ and a model $f(x;\varphi)$, where $x \in \mathbb{R}^{N_x}$ is the input vector and $\varphi \in \mathbb{R}^{N_\varphi}$ is the parameter vector, our goal is to find parameters φ such that $f(x;\varphi)$ is a good approximation to the observations, $\mathcal{D}$, as well as future predictions. The distribution of the parameters conditioned on the data is given by the Bayes theorem,

$$p(\varphi|\mathcal{D}) = \frac{p(\mathcal{D}|\varphi)\; p(\varphi)}{p(\mathcal{D})} \; . \tag{1}$$

The probability distribution $p(\varphi)$ is called the prior distribution and it encodes all the available information on the parameters prior to observing any data. The denominator, $p(\mathcal{D})$, is usually referred to as the evidence of the data and used for model selection [3]. Under the assumption that the data are independent and normally distributed around the output of the model,

$$d_i = f(x_i;\varphi) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_n) \; , \tag{2}$$

the likelihood function $p(\mathcal{D}|\vartheta)$ takes the form,

$$p(\mathcal{D}|\vartheta) = \mathcal{N}(\mathcal{D} \,|\, F(X,\varphi), \sigma_n I) \; , \tag{3}$$

where $\vartheta = (\varphi^\top, \sigma_n)^\top$ is the parameter vector that contains both, the model and the noise, parameters and $F(X,\varphi) = (f(x_1;\varphi), \ldots, f(x_{N_\mathcal{D}};\varphi))$. The assumption that the observations are independent is used here in order to simplify the presentation and correlations between the observations can be included without changing the general formulation presented here.

2.2 Transitional Markov Chain Monte Carlo (TMCMC)

TMCMC is a population based algorithm with an annealing scheme: sampling from a sequence of intermediate distributions,

$$p_j(\vartheta) \propto p(\mathcal{D}|\vartheta)^{\zeta_j} p(\vartheta) \; , \tag{4}$$

for $j = 1, \ldots, m$ and $0 = \zeta_1 < \ldots < \zeta_m = 1$, that converges to the posterior distribution $p(\vartheta|\mathcal{D}) \propto p(\mathcal{D}|\vartheta)p(\vartheta)$ when $\zeta_m = 1$.

The algorithm at the first stage draws N_1 samples from the prior distribution and at the $j+1$ stage uses N_j samples from the distribution p_j to obtain N_{j+1} samples from the distribution p_{j+1}. Let $\Theta_j = \{\vartheta_{j,k} | k = 1, \ldots, N_j\}$ be the samples obtained at the j-th step from p_j. The following procedure gives samples from p_{j+1}:

1. Draw N_{j+1} samples from the set Θ_j with probability of the sample $\vartheta_{j,k}$ to be selected equal to

$$\hat{w}_{j,k} = \frac{w_{j,k}}{\sum_{k=1}^{N_i} w_{j,k}} , \tag{5}$$

where $w_{j,k} = p(\mathcal{D}|\vartheta_{j,k})^{\zeta_{j+1}-\zeta_j}$. Put the new samples in the set $\tilde{\Theta}_{j+1}$ and set

$$S_j = \frac{1}{N_j} \sum_{k=1}^{N_j} w_{j,k} . \tag{6}$$

2. For each sample in $\tilde{\Theta}_{j+1}$ perform TMCMC with Gaussian proposal distribution and covariance matrix $\beta^2 \Sigma_s^{(j)}$. Here, β is a scaling parameter and $\Sigma_s^{(j)}$ is the sample covariance at the j-th stage given by,

$$\Sigma_s^{(j)} = \sum_{k=1}^{N_j} \hat{w}_{j,k} (\vartheta_{j,k} - \bar{\vartheta}_j)(\vartheta_{j,k} - \bar{\vartheta}_j)^\top , \tag{7}$$

where $\bar{\vartheta}_j = \sum_{k=1}^{N_j} \hat{w}_{j,k} \vartheta_{j,k}$. Set the chain length equal to a predefined parameter ℓ_{max}.

The pseudocode of the algorithm is presented in Algorithm 1. The authors in [25] suggested that the MCMC step described in lines 11-14 of Algorithm 1 results in a bias accumulated in each stage. In the original TMCMC the MCMC is performed for each *unique* sample in $\tilde{\Theta}_{j+1}$ with chain length equal to number of occurrences of the sample. In [25] it is shown that in order for the bias to be alleviated all samples in $\tilde{\Theta}_{j+1}$ should perform an MCMC step with chain length equal to a parameter ℓ_{max}, which is usually set equal to 1. The improved algorithm is called BASIS and an efficient implementation can be found in the Π4U framework [12].

An advantageous property of the TMCMC algorithm is that it can be efficiently parallelized since the computationally expensive part of the algorithm, which is the MCMC step that involves the likelihood evaluation, is independent for each sample. A key computational benefit of the BASIS algorithm is that all MCMC chains have equal length and thus the work load is balanced among the processors.

Moreover, an important byproduct of the algorithm is that the evidence of the data is approximated by

$$p(\mathcal{D}) \approx \prod_{j=1}^{m-1} S_j , \tag{8}$$

where m is the total number of stages in the TMCMC algorithm. A proof of this can be found in [8]. The idea of approximating the evidence of the data by (8) can also be found in [7, 16] in the context of thermodynamic integration. The main difference between these algorithms and TMCMC is that in TMCMC, the annealing schedule, $\{\zeta_j | j = 1, \ldots, m\}$, is estimated adaptively according to the scheme described in line 6 of Algorithm 1. In the thermodynamic integration the annealing schedule is chosen a priori.

Finally, we note that the parameter β^2 was proposed in [8] to be set equal to 0.04. In [5] the authors propose an adaptive choise of β^2 such that a predefined acceptance rate is achieved. We return to the issue of choosing β^2 in Section 3.2. For the rest of the paper, in order to be consistent with the notation introduced later, we will denote β^2 by ε.

2.3 Manifold Metropolis Adjusted Langevin Algorithms

Let $\pi : \Theta \to \mathbb{R}$ be a probability distribution function where $\Theta \subset \mathbb{R}^{N_\vartheta}$. The Metropolis-Hastings (MH) algorithm is a Markov Chain Monte Carlo (MCMC) method for obtaining samples from π according to the following iterative scheme, starting from sample ϑ_0:

1. propose a sample $\vartheta^\star$ according to a probability function $q(\vartheta^\star|\vartheta_j)$,

Algorithm 1 BASIS (TMCMC)

Input: Likelihood function $p(\mathcal{D}|\vartheta)$, prior distribution $p(\vartheta)$
N_j, N_{max} – number of samples per stage, maximum number of stages
β, γ – threshold parameter, scale parameter

Output: Θ_{final} – a set of samples from $p(\vartheta|\mathcal{D})$
S – estimation for the evidence $p(\mathcal{D})$

Draw initial sample set $\Theta_1 = \{\vartheta_{1,k}|k = 1, \ldots, N_1\}$ from prior
Initialize $j \leftarrow 1$, $\zeta_1 \leftarrow 0$, $S \leftarrow 1$
repeat
  Choose ζ_{j+1} such that CV of $w_{j,k} < \gamma$ and $\zeta_{j+1} \leq 1$
  Calculate $w_{j,k}$ with the chosen ζ_{j+1}
  $S \leftarrow S \cdot \frac{1}{N_j} \sum_{j=1}^{N_j} w_{j,k}$
  Obtain $\tilde{\Theta}_{j+1}$ by drawing N_{j+1} samples from the set Θ_j with probabilities $\propto w_{j,k}$
  Set Σ the weighted covariance given by (7)
  for each sample in $\tilde{\Theta}_{j+1}$ **do**
    Perform MCMC with length equal to ℓ_{max} and proposal distribution $q(\cdot|\vartheta) = \mathcal{N}(\cdot|\vartheta, \beta^2\Sigma)$
    Add resulting samples in Θ_{j+1}
  end for
  $j \leftarrow j + 1$
until $\zeta_j = 1$ or $j > N_{max}$
$\Theta_{final} \leftarrow \Theta_j$

2. set $\vartheta_{j+1} = \vartheta^\star$ with probability

$$\alpha(\vartheta^\star, \vartheta_j) = \frac{\pi(\vartheta^\star) q(\vartheta_j|\vartheta^\star)}{\pi(\vartheta_j) q(\vartheta^\star|\vartheta_j)} , \tag{9}$$

and $\vartheta_{j+1} = \vartheta_j$ with probability $1 - \alpha$.

In the limit the samples obtained using the MH algorithm will be distributed according to π. The proposal distribution q is usually chosen to be a Gaussian distribution centered at ϑ_j with covariance matrix $\sigma^2 I$, i.e., $q(\cdot|\vartheta_j) = \mathcal{N}(\cdot|\vartheta_j, \sigma^2 I)$ and I the identity matrix in $\mathbb{R}^{N_\vartheta}$. The parameter σ needs to be tuned depending on the probability p; if σ is too small the proposals will be local and the chain will not be able to efficiently explore the parameter space, if σ is too big there will many rejected samples leading to slow convergence.

An improved proposal scheme is based on the observation that the random variable that satisfies the SDE,

$$d\vartheta_t = \frac{1}{2} \nabla \log \pi(\vartheta)\, dt + dW_t , \tag{10}$$

where W_t an N_ϑ dimensional Wiener process, has π as stationary distribution. Then, the Euler-Maruyama discretization is given by

$$\vartheta_{n+1} = \vartheta_n + \frac{\varepsilon}{2} \nabla \log \pi(\vartheta_n) + \sqrt{\varepsilon} W_n, \quad W_n \sim \mathcal{N}(0, I) , \tag{11}$$

where ε is the time step of the discretization. Since ε introduces error, π is not anymore the equilibrium distribution of ϑ_t and thus cannot be sampled directly by solving (10). Instead, the sample ϑ_{n+1} is used as a proposal in the MH algorithm. In other words, the proposal distribution in the MH algorithm is given by

$$q(\cdot|\vartheta_n) = \mathcal{N}(\cdot \,|\, \vartheta_n + \frac{\varepsilon}{2} \nabla \log \pi(\vartheta_n), \varepsilon I) . \tag{12}$$

This scheme is known as the Metropolis Adjusted Langevin Algorithm (MALA). Notice that the original MH algorithm with Gaussian proposal distribution corresponds to the MALA algorithm for the SDE

$$d\vartheta_t = dW_t , \tag{13}$$

which corresponds to an isotropic diffusion in $\mathbb{R}^{N_\vartheta}$ and is usually referred to as the random walk MH algorithm (RWMH).

Although the proposals based on (10) follow the direction of the maximum change, guiding ϑ to regions of high probability, the isotropic diffusion may be inappropriate in the presence of highly correlated random variables. The following SDE offers a better proposal scheme,

$$d\vartheta_t = \frac{1}{2}\Sigma\nabla \log \pi(\vartheta)\, dt + \sqrt{\Sigma} dW_t \ , \tag{14}$$

where the correlation of the variables is encoded in the constant and positive definite matrix Σ. The algorithm is known as the pre-conditioned MALA and was proposed in [21].

A variation of the pre-conditioned MALA was proposed in [10] and is based on following SDE with position dependent covariance matrix

$$d\vartheta_t = \frac{1}{2}G^{-1}(\vartheta_t)\nabla \log \pi(\vartheta_t)dt + \Omega(\vartheta_t)dt + G^{-\frac{1}{2}}(\vartheta_t)dW_t, \tag{15}$$

where

$$\Omega_i(\vartheta_t) = |G(\vartheta_t)|^{-\frac{1}{2}} \sum_j \frac{\partial}{\partial \vartheta_j}\left[G^{-1}_{i,j}(\vartheta_t)|G(\vartheta_t)|^{\frac{1}{2}}\right] \ , \tag{16}$$

and G is some positive definite matrix. Due to the fact that (15) describes a diffusion in a manifold, defined in local coordinates by G [10], the resulting MCMC algorithm is called manifold MALA (mMALA). In [26] the authors show that the correct form of Ω should be

$$\Omega_i(\vartheta_t) = \frac{1}{2}\sum_j \frac{\partial}{\partial \vartheta_j} G^{-1}_{i,j}(\vartheta_t) = -\frac{1}{2}\sum_j \left[G^{-1}(\vartheta_t)\frac{\partial}{\partial \vartheta_j} G_{i,j}(\vartheta_t) G^{-1}(\vartheta_t)\right]_{i,j} \ , \tag{17}$$

and prove that (16) and (17) describe equal diffusions under the condition $\partial_{x_j} G_{k,m}(\vartheta) = \partial_{x_k} G_{j,m}(\vartheta)$. The authors refer to the resulting sampling scheme as position-depended MALA (pMALA). A simplified version, under the assumption that the manifold has constant curvature, can be obtained by setting $\Omega = 0$ leading to proposals,

$$q(\cdot|\vartheta) = \mathcal{N}(\cdot \,|\, \vartheta + \frac{\varepsilon}{2} G^{-1}(\vartheta)\, \nabla \log \pi(\vartheta), \varepsilon G^{-1}(\vartheta)) \ . \tag{18}$$

The resulting scheme is known as simplified manifold MALA (smMALA) [10].

There are two remaining issues: (a) what is the optimal choice for G and (b) how the parameter ε should be chosen. Regarding the second question, in [21] optimal scaling with respect to the dimension of ϑ for various MH schemes is derived, MALA included, and in [6] optimal scaling for a class of mMALA algorithms is provided for a wide class of distributions. We return to this issue in Section 3.2 where we discuss a heuristic procedure for automatic tuning of ε in the framework of TMCMC.

In order to answer the second question we consider a specific form for π as a posterior of Bayesian inference problem, i.e., $\pi(\vartheta) = p(\vartheta|\mathcal{D}) \propto p(\mathcal{D}, \vartheta) = p(\mathcal{D}|\vartheta)p(\vartheta)$. Then there are two widely used choises: the negative of the Hessian of $\log p(\mathcal{D}, \vartheta)$,

$$\begin{aligned} G(\vartheta) = \mathcal{H}(\vartheta) &:= -\nabla^2 \log p(\mathcal{D}, \vartheta) \\ &= -\nabla^2 \log p(\mathcal{D}|\vartheta) - \nabla^2 \log p(\vartheta) \ , \end{aligned} \tag{19}$$

and the Fisher information of $p(\mathcal{D}, \vartheta)$,

$$\begin{aligned} G(\vartheta) = \mathcal{I}(\vartheta) &:= -\mathbb{E}_{\mathcal{D}|\vartheta}\left[\nabla^2 \log p(\mathcal{D}, \vartheta)\right] \\ &= -\mathbb{E}_{\mathcal{D}|\vartheta}\left[\nabla^2 \log p(\mathcal{D}|\vartheta)\right] - \nabla^2 \log p(\vartheta) \ , \end{aligned} \tag{20}$$

which is the Fisher information of the likelihood function minus the Hessian of the prior distribution. In [17] an alternative derivation of the mMALA with the Hessian as covariance matrix is presented where a connection with optimization methods is made, hence they call their method Stochastic Newton MCMC. In Section 3.1 we discuss the most suitable choice of G.

3 Manifold Transitional Markov Chain Monte Carlo

In this section we propose a modification of the original TMCMC algorithm. In the TMCMC algorithm the proposal distribution at the j-th stage at point ϑ is a Gaussian density function

centred at ϑ with a constant covariance matrix, i.e., $\vartheta' \sim \mathcal{N}(\vartheta, \beta^2\Sigma_s^{(j)})$, where $\Sigma_s^{(j)}$ is the weighted sample covariance matrix at stage j given by (7). Here, we intend to improve the quality of TMCMC samples by using a proposal scheme based on the manifold MALA algorithms discussed in Section 2.3.

In the j-th stage of TMCMC samples from the distribution

$$p_j(\vartheta) = \frac{p(\mathcal{D}|\vartheta)^{\gamma_j} p(\vartheta)}{p_j(\mathcal{D})} \ , \tag{21}$$

must be collected. In order to use the proposal scheme (15) the gradient of $\log p_j$ should be computed,

$$\nabla \log p_j(\vartheta) = \gamma_j \nabla \log \ p(\mathcal{D}|\vartheta) + \nabla \log p(\vartheta) \ . \tag{22}$$

For the rest of the presentation we restrict the discussion in uniform prior distributions, thus the last term in (22), as well as the last term in $G(\vartheta)$ in (19) and (20), vanishes. In this case the diffusion (15) is written as,

$$d\vartheta_t = \frac{\gamma}{2}\Sigma_\gamma \nabla \log p(\mathcal{D}|\vartheta_t)dt + \Omega_\gamma(\vartheta_t)dt + \sqrt{\Sigma_\gamma(\vartheta_t)}dW_t \ , \tag{23}$$

where $\Sigma_\gamma = G_\gamma^{-1}$ with G_γ the metric corresponding to $p^\gamma(\mathcal{D}|\vartheta)$ and Ω_γ is given by (17) with G substituted by G_γ. Notice, that under the modelling assumption (3) and with G given by (19) and (20) it can be shown that,

$$\Sigma_\gamma(\vartheta) = \gamma^{-1}\Sigma(\vartheta) \quad \text{ and } \quad \Omega_\gamma(\vartheta) = \gamma^{-1}\Omega(\vartheta) \ . \tag{24}$$

Moreover, the gradient of the log-likelihood function is given by,

$$\frac{\partial}{\partial \vartheta_k} \log p(\mathcal{D}|\vartheta) = \begin{cases} \sigma_n^{-2} \sum_{i=1}^{N_d} \left(d_i - f_i\right) \frac{\partial}{\partial \varphi_k} f_i, & k = 1, \dots, N_\vartheta - 1 \ , \\ -N_d \sigma_n^{-1} + \sigma_n^{-3} \sum_{i=1}^{N_d} \left(d_i - f_i\right)^2, & k = N_\vartheta \ , \end{cases} \tag{25}$$

where $f_i = f(x_i; \varphi)$ and the Hessian is given by,

$$\frac{\partial^2}{\partial \vartheta_k \partial \vartheta_\ell} \log p(\mathcal{D}|\vartheta) = \begin{cases} \sigma_n^{-2} \sum_{i=1}^{N_d} \left(d_i - f_i\right) \left(\frac{\partial}{\partial \varphi_k \varphi_\ell} f_i - \frac{\partial}{\partial \varphi_\ell} f_i \frac{\partial}{\partial \varphi_k} f_i \right), & k, \ell = 1, \dots, N_\vartheta - 1 \ , \\ -2\sigma_n^{-3} \sum_{i=1}^{N_d} \left(d_i - f_i\right) \frac{\partial}{\partial \varphi_k} f_i, & k = 1, \dots, N_\vartheta - 1, \ \ell = N_\theta \ , \\ -N_d - 3\sigma_n^{-4} \sum_{i=1}^{N_d} \left(d_i - f_i\right)^2, & k, \ell = N_\vartheta \ . \end{cases} \tag{26}$$

We note that the computation of the gradient and the Hessian of $p(\mathcal{D}|\vartheta)$ involves computation of the the derivatives of the observable function f. We discuss the details of this computation in Section 4.

3.1 Choices and corrections for the covariance matrix

Here we discuss two possible choices of the metric G, the Hessian and the Fisher information, defined in (19) and (20), respectively. There are three potential problems with the matrix G:

(a) G is not invertible,

(b) G is invertible with some negative eigenvalues, thus G^{-1} is negative semi-definite,

(c) G is invertible and G^{-1} positive semi-definite and some eigenvalues of G are very small, respectively some eigenvalues of G^{-1} are very large, leading to a proposal Gaussian distribution that has a large volume outside the bounds of the prior distribution.

The first case is proposed to be fixed by setting $G = \Sigma_s^{(j)}$, the sample covariance matrix at the j-th stage of TMCMC. The second problem can appear only when $G = \mathcal{H}$ since the Fisher information is always positive semi-definite. The presence of zero eigenvalues is treated in the third case. One possible solution to this is to set $G = \Sigma_s^{(j)}$ as in the non-invertible case. As one might expect, this solution is far from optimal since all the information coming from directions with positive eigenvalues is discarded. In fact, it is usually the case that only a few eigendirections appear with negative eigenvalues and the rest with positive. In order to preserve this information, we choose to substitute the negative eigenvalues of G^{-1} with a predefined positive number and this number is chosen to be equal to the smallest eigenvalue of Σ_p. Other proposals for fixing the negative definiteness of the covariance matrix include the SoftAbs method presented in [4] and a modified Choleksy decomposition used in [15].

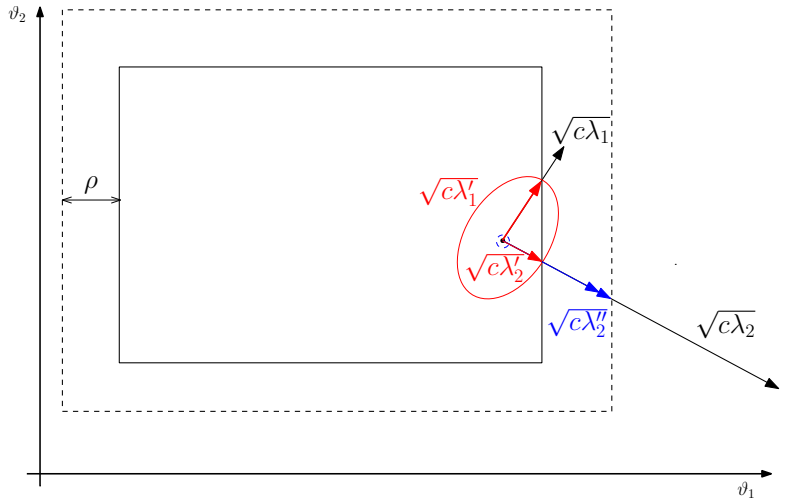

Figure 1: Illustration of the eigenvalue adaptation in order to treat unidentifiable directions. Here, c is equal to the inverse of the cumulative χ^2_k distribution with k degrees of freedom, $c = \chi_k^{-2}(\alpha)$. Note that instead of adapting λ_2 to the boundaries of the prior distribution (continuous line) we adapt to an extended boundary (dashed line) which is defined as an extension of $100\rho\%$ of the length of the prior boundaries.

The third problem appears in the case of unidentifiable parameters or parameter combinations. Along these directions the likelihood of the data stays constant and thus the second derivative is zero. In many practical situations, the issue of *computational unidentifiability* appears [18] in which some parameters are identifiable in a small region of the probability space but close to unidentifiable in a large region of the space. Unidentifiable directions correspond to zero, or close to zero, eigenvalues in the Hessian or the Fisher information, leading to large eigendirections in the proposal Gaussian distribution.

We proceed to treat the problems arising from the large eigenvalues as follows: Let $\Sigma = Q\,\Lambda\,Q^\top \in \mathbb{R}^{k,k}$ be the eigendecomposition of the covariance matrix of a normal distribution centered at the origin, with $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_k)$, λ_i the eigenvalues of Σ and Q a square matrix whose columns q_i are the eigenvectors corresponding to the λ_i. Then α percentage of the probability mass lies inside the ellipsoid described by the equation

$$(Q^\top x)^\top\,\Lambda^{-1}\,(Q^\top x) = \chi_k^{-2}(\alpha)\;, \tag{27}$$

where χ_k^{-2} is the inverse of the cumulative χ^2_k distribution with k degrees of freedom [24]. The semi-axes of the ellipsoid (27) are given by $\sqrt{\lambda_i \chi_k^{-2}(\alpha)}q_i$. Then we adapt the eigenvalues of Σ by finding constants c_i such that $\pm\sqrt{\widehat{\lambda}_i \chi_k^{-2}(\alpha)}q_i$ lies inside the bounds of the prior distribution, with $\widehat{\lambda}_i = c_i\lambda_i$. We have observed that this approach leads to small eigenvalues near the boundaries and thus the proposal distribution may be concentrated near the boundaries. We propose to resolve this issue by adapting $\pm\sqrt{\lambda_i \chi_k^{-2}(\alpha)}$ to an extended boundary that is defined as a percentage, ρ, of the length of the boundaries of the prior probability. We will denote the corrected covariance by $\widehat{\Sigma} = Q\widehat{\Lambda}Q^\top$ and $\widehat{\Lambda} = C\Lambda$ with $C = \mathrm{diag}(c_1, \ldots, c_{N_\vartheta})$. For an illustration of this approach see Figure 1. The advantages of the extended boundaries approach, as well as the selection of the parameter ρ, is presented in a truncated multivariate Gaussian in Section 5.1.

$$\begin{array}{ccccccc} G & \longrightarrow & \Sigma = G^{-1} & \longrightarrow & \widehat{\Sigma} = \widehat{G}^{-1} & \longrightarrow & \widehat{G} \\ \| & & \| & & \| & & \| \\ Q\Lambda Q^\top & & Q\Lambda^{-1}Q^\top & & Q\widehat{\Lambda}^{-1}Q^\top & & Q\widehat{\Lambda}Q^\top \end{array}$$

Figure 2: The steps of correcting the Hessian or Fisher information to $\widehat{G}$.

3.2 Sampling schemes

In this section we summarize the proposal schemes we use in the TMCMC algorithm. For the proposal distribution we will use the unified notation,

$$q(\vartheta^\star|\vartheta) = \mathcal{N}\big(\vartheta^\star \mid D(\vartheta;\varepsilon,\gamma), C(\vartheta;\varepsilon,\gamma)\big) \ , \tag{28}$$

with D and C the mean and the covariance matrix of the Gaussian proposal distribution, repsectively. In the original TMCMC we have $D(\vartheta;\varepsilon,\gamma) = \vartheta$ and $C(\vartheta;\varepsilon,\gamma) = \varepsilon\Sigma_s^{(j)}$. Since $\Sigma_s^{(j)}$ is a constant matrix the proposal distribution is symmetric and thus the acceptance ratio is independent of q.

The simplified manifold TMCMC (smTMCMC) (18) is obtained by assuming that the metric, G_γ, is locally constant. This assumption leads to $\Sigma_\gamma(\vartheta) = \Sigma_\gamma(\vartheta^\star)$ and $\Omega_\gamma(\vartheta) = 0$. The covariance matrix $\Sigma_\gamma(\vartheta)$ is substituted by $\widehat{\Sigma}_\gamma(\vartheta)$, the corrected covariance according to the scheme presented in Section 3.2. Finally, the mean and the variance of the proposal distribution are given by $D(\vartheta;\varepsilon,\gamma) = \vartheta + \frac{\varepsilon\gamma}{2}\widehat{\Sigma}_\gamma(\vartheta)\nabla \log p(\mathcal{D}|\vartheta)$ and $C(\vartheta;\varepsilon,\gamma) = \varepsilon\Sigma_\gamma(\vartheta)$, respectively. Notice that the proposal distribution is not symmetric and thus must be included in the calculation of the acceptance ratio, see Table 1.

If we do not make the assumption of locally constant metric then Ω_γ is not necessarily zero. After correcting the covariance matrix, Σ_γ, the Ω term is written as,

$$\widehat{\Omega}_{\gamma,i}(\vartheta) = -\frac{1}{2}\sum_j \Big[\widehat{\Sigma}_\gamma(\vartheta)\ \frac{\partial}{\partial\vartheta_j}\widehat{G}_{\gamma,i,j}(\vartheta)\ \widehat{\Sigma}_\gamma(\vartheta)\Big]_{i,j} \ , \tag{29}$$

with $\widehat{G}_\gamma = \widehat{\Sigma}_\gamma^{-1}$. In order to express the unknown derivative of $\widehat{G}$ in terms of the known derivative of G we follow the same procedure as in [4], see also Appendix A,

$$\frac{\partial\widehat{G}_\gamma}{\partial\vartheta_j} = Q\Big(J\circ\big(Q^\top\frac{\partial G_\gamma}{\partial\vartheta_j}Q\big)\Big)Q^\top \ , \tag{30}$$

where $\circ$ denotes the Hadamard product and

$$J_{i,j} = \begin{cases} \frac{\widehat{\lambda}_i-\widehat{\lambda}_j}{\lambda_i-\lambda_j} = \frac{c_i\lambda_i-c_j\lambda_j}{\lambda_i-\lambda_j}, & i\neq j \\ \frac{\partial\widehat{\lambda}_i}{\partial\lambda_i} = c_i, & i=j \ . \end{cases} \tag{31}$$

We call the resulting algorithm position dependend TMCMC (pTMCMC).

Tuning the scale parameter. Finally, in all the sampling algorithms the parameter ε must be tuned appropriately. A usual practice is to vary ε until the acceptance rate of the algorithm hits a target acceptance rate. Following the results in [23] we set the target acceptance rate for TMCMC, which is based in random walk MH, equal to 0.234 and for the Langevin diffusion based TMCMC equal to 0.574. Moreover it can be shown that ε scales like N_ϑ^{-1} for RWMH and like $N_\vartheta^{-1/3}$ for MALA [23] as $N_\vartheta\to\infty$. One way to find the optimal scaling is to run some preparatory, relatively small, simulations with various ε and choose the one that gives the optimal acceptance rate.

An adaptive method to find the optimal scaling is proposed in [5] in the framework of TMCMC. In this method the sample sets at the j-th stage of TMCMC are divided into subsets. These subsets are being run sequentially and information of the mean acceptance rate is being passed to the next subset. The scaling parameter is being tuned up or down depending on how far is the mean acceptance rate from the target acceptance rate.

Choosing the metric. In the aforementioned sampling schemes either the Hessian or the Fisher information can be used as the underlying metric. Notice that the Hessian needs the computation of second order derivatives while the Fisher information first order derivatives. Moreover, since Fisher information is always positive semidefinite, we expect it to perform better than the Hessian. However, in some cases the Fisher information is not explicitly known, e.g., in the case of Gaussian mixtures, and thus the use of Hessian is inevitable.

Finally, in the pTMCMC scheme, where the computation of $\widehat{\Omega}$ is involved, the order of the derivatives needed is increased by one. Thus, for $G = \mathcal{H}$ and $G = \mathcal{I}$, third and second order derivatives, respectively, should be computed. In order to avoid the computation of third order derivatives we choose not to implement the Hessian in these two sampling schemes.

The performance of the algorithms presented here is discussed in detail in Section 5. In Table 1 we present a summary of the proposal distribution and the resulting acceptance ratio for the various sampling schemes.

name	**mean**, $D(\vartheta;\varepsilon,\gamma)$	**covariance**, $C(\vartheta;\varepsilon,\gamma)$	**acceptance ratio**, $\alpha(\vartheta^\star,\vartheta)$
TMCMC	ϑ	$\varepsilon\Sigma_p$	$\dfrac{p(\mathcal{D}\|\vartheta^\star)}{p(\mathcal{D}\|\vartheta)}$
smTMCMC	$\vartheta + \dfrac{\varepsilon\gamma}{2}\widehat{\Sigma}_\gamma(\vartheta)\nabla \log p(\mathcal{D}\|\vartheta)$	$\varepsilon\widehat{\Sigma}_\gamma(\vartheta)$	$\dfrac{p(\mathcal{D}\|\vartheta^\star)\,\mathcal{N}\big(\vartheta \,\|\, D(\vartheta^\star), C(\vartheta)\big)}{p(\mathcal{D}\|\vartheta)\,\mathcal{N}\big(\vartheta^\star \,\|\, D(\vartheta), C(\vartheta)\big)}$
pTMCMC	$\vartheta + \dfrac{\varepsilon\gamma}{2}\widehat{\Sigma}_\gamma(\vartheta)\nabla \log p(\mathcal{D}\|\vartheta) + \varepsilon\widehat{\Omega}_\gamma(\vartheta)$	$\varepsilon\widehat{\Sigma}_\gamma(\vartheta)$	$\dfrac{p(\mathcal{D}\|\vartheta^\star)\,\mathcal{N}\big(\vartheta \,\|\, D(\vartheta^\star), C(\vartheta^\star)\big)}{p(\mathcal{D}\|\vartheta)\,\mathcal{N}\big(\vartheta^\star \,\|\, D(\vartheta), C(\vartheta)\big)}$

Table 1: Proposal distribution, $q(\vartheta^\star|\vartheta) = \mathcal{N}\big(\vartheta^\star \,|\, D(\vartheta;\varepsilon,\gamma), C(\vartheta;\varepsilon,\gamma)\big)$, and acceptance ratio, $\alpha(\vartheta^\star,\vartheta)$, for the TMCMC algorithms described in Section 3.2.

4 Application to Ordinary Differential Equations

In this section we show how the derivatives of the model output, f, with respect to the model parameters, φ, can be computed. These quantities appear in the evaluation of the the derivative of the log-likelihood function with respect to the parameters, see (25) and (26).

Let $f := f(t;\varphi)$ be an observable function on the solution of the ODE system,

$$\begin{cases} \dot{Y}(t) = G(Y;\varphi), \quad t \in [0,T] \\ Y(0) = g(\varphi)\ , \end{cases} \tag{32}$$

where $Y \in \mathbb{R}^{N_Y}$ and $\varphi \in \mathbb{R}^{N_\varphi}$ a the vector of parameters. Here, we assume that G is a smooth function such that all the derivatives used bellow are well defined. In order to find the first and second order derivatives, needed for the manifold algorithms the following extended system must be solved,

$$\begin{cases} \dot{Y} = G(Y;\varphi) \\ \dot{S}_k = G^1_k(Y,S;\varphi), \quad k = 1,\dots,N_\varphi \\ \dot{H}_{k,\ell} = G^2_{k,\ell}(Y,S,H;\varphi), \quad k,\ell = 1,\dots,N_\varphi \\ Y(0) = g(\varphi), \;\; S(0) = g^1_k(\varphi), \;\; H(0) = g^2_{k,\ell}(\varphi)\ , \end{cases} \tag{33}$$

where G^1_k and $G^2_{k,\ell}$ is the first total derivative of G with respect to φ_k and second total derivative of G with respect to φ_k and φ_ℓ, respectively, and $g^1_k = \frac{\partial}{\partial\varphi_k} g$ and $g^2_{k,\ell} = \frac{\partial^2}{\partial\varphi_k\partial\varphi_\ell} g$.

The function $G^1_k = D_{\varphi_k} G$ for $k = 1,\dots,N_\varphi$ is given by,

$$G^1_k = AS_k + B_k\ , \tag{34}$$

where

$$A_{ij} := A_{ij}(Y;\varphi) = \frac{\partial}{\partial Y_j} G_i(Y;\varphi)\ , \tag{35}$$

for $i,j = 1,\dots,N_Y$ and

$$B_k := B_k(Y;\varphi) = \frac{\partial}{\partial\varphi_k} G(Y;\varphi)\ . \tag{36}$$

The function $G^2_{k,\ell} = D_{\varphi_k} D_{\varphi_\ell} G$ for $k,\ell = 1,\dots,N_\varphi$ is given by,

$$G^2_{k,\ell} = AH_{k,\ell} + (I \otimes S_k^\top) C_{k,\ell} (\mathbb{1} \otimes S_\ell) + D_k S_\ell + D_\ell S_k + J_{k,\ell}\ , \tag{37}$$

where A is defined in (35), $\otimes$ is the Kronecker product, $I \in \mathbb{R}^{N_Y \times N_Y}$ the identity matrix and $\mathbb{1} = (1,\dots,1) \in \mathbb{R}^{N_Y}$. The matrix C is a block diagonal matrix $C_{k,\ell} = \mathrm{diag}(C_{k,\ell,1},\dots,C_{k,\ell,N_Y})$ with the block matrices given by

$$C_{k,\ell,i} := C_{k,\ell,i}(Y;\varphi) = \frac{\partial^2}{\partial Y_k \partial Y_\ell} G_i(Y;\varphi)\ . \tag{38}$$

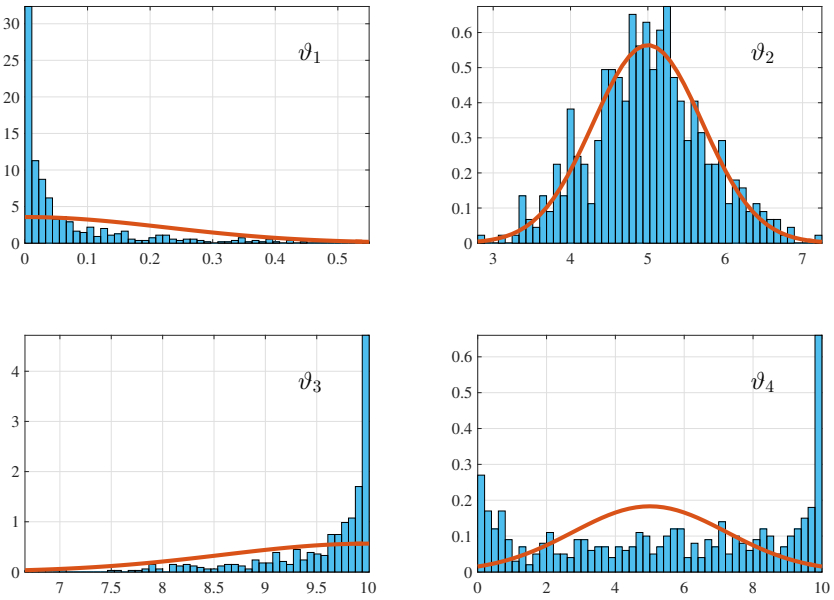

Figure 3: Marginal histograms of the truncated Gaussian distribution (41) compared with the exact distribution (red line). The sampling algorithm is the smMALA algorithm with no extended boundaries. Notice that the sampling algorithm incorrectly concentrates samples near the boundaries of the prior.

The matrix D_k and the vector $J_{k,\ell}$ are given by

$$D_{k,i,j} := D_{k,i,j}(Y;\varphi) = \frac{\partial^2}{\partial\varphi_k \partial Y_j} G_i(Y;\varphi)\ , \tag{39}$$

$$J_{k,\ell} := J_{k,\ell}(Y;\varphi) = \frac{\partial^2}{\partial\varphi_k \partial\varphi_\ell} G(Y;\varphi)\ , \tag{40}$$

for $i, j = 1, \ldots, N_Y$.

A Matlab function is provided which given functions G and g, by performing symbolic calculations, gives as output the functions G^1_k and $G^2_{k,\ell}$ as Matlab function handles, which can be downloaded from [1].

5 Numerical examples

In this section we demonstrate the performance improvements achieved by combining TMCMC and Langevin diffusions. In the first example we test our algorithm in a truncated Gaussian distribution. This test demonstrates the need of the extended boundary approach discussed in Section 3.1. In the second example the scaling of the error of smTMCMC is compared to that of TMCMC in a multivariate Gaussian distribution. Finally, in the third example we sample the posterior distribution of a Bayesian inference problem of a 8 dimensional pharmacodynamics model. The challenging part of sampling this particular distribution is that it has large manifolds of unidentifiable or nearly unidentifiable parameters. In this example the TMCMC algorithm completely fails to explore the parameter space while the smTMCMC algorithm provides good quality samples.

In all the numerical tests, where Fisher information was used and thus the application of pTMCMC was possible, the pTMCMC algorithm was also implemented. However the difference between smTMCMC and pTMCMC is insignificant and in most cases the results were indistinguishable.

5.1 Truncated Gaussian distribution

In this section we illustrate the ability of the proposed algorithm to sample correctly distributions with mass concentrated at the boundaries of the prior algorithm. The strategy of adapting the

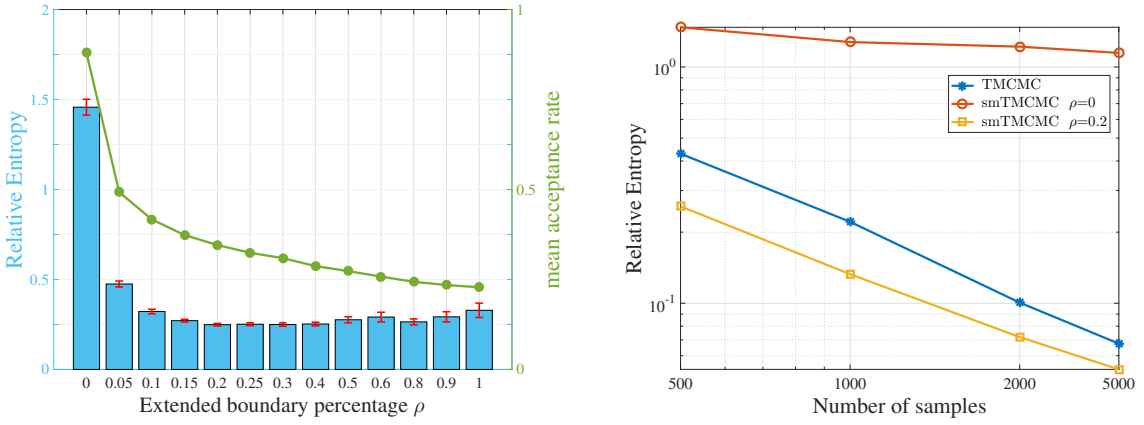

(a) (b)

Figure 4: (a) Relative entropy (KL divergence) between the estimated and the truncated Gaussian distribution of (41) as a function of the parameter ρ for the smTMCMC algorithm. (b) The same quantity as a function of the sample size of the sampling algorithm for TMCMC, smTMCMC with $\rho = 0, 0.2, 0.5, 1$.

eigenvalues of the covariance matrix of the proposal distribution by extending the boundaries of the prior distribution, as discussed in Section 3.1, is compared with the default approach. In the default approach the eigenvalues are adapted directly to the boundaries of the prior. Thus, the corrected covariance matrix, $\widehat{\Sigma}$, at samples near the boundaries will have some very small eigenvalues and the MCMC chain starting at these points will be locally trapped.

We now consider the truncated Gaussian probability distribution

$$p(\vartheta) = \prod_{i=1}^{4} \mathcal{N}(\vartheta_i \,|\, \mu_i, \sigma_i^2)\, \mathcal{U}(\vartheta_i \,|\, 0, 10) \,, \tag{41}$$

with $\mu = (0, 5, 10, 9)$ and $\sigma^2 = (0.05, 0.5, 2, 5)$. The marginal distribution of ϑ_i is depicted with a red line in Figure 3. The histograms in Figure 3 are obtained with the smMALA algorithm and G the Fisher information of (41) using 500 samples and $\varepsilon = 1$. It is evident from this example that the samples are concentrated at the boundaries of the prior distribution where there is significant mass of the the Gaussian distribution. Notice that the variable ϑ_2, which has probability away from the boundaries, is sampled correctly.

Next, we sample (41) using smMALA with extended boundaries and we study the effect of the the parameter ρ to the accuracy of the sampling. We measure the accuracy using the relative entropy or Kullback-Leibler (KL) divergence of p from $\tilde{p}$,

$$D_{KL}(\tilde{p} \,\|\, p) = \int_{-\infty}^{\infty} \tilde{p}(\vartheta) \log \frac{\tilde{p}(\vartheta)}{p(\vartheta)}\, d\vartheta \,, \tag{42}$$

with p absolutely continuous with respect to $\tilde{p}$, i.e., $p(\vartheta) = 0$ implies $\tilde{p}(\vartheta) = 0$. Here we take $\tilde{p}$ to be the estimated probability distribution using a sampling algorithm and p the target distribution (41). The KL divergence is not a metric, since it is not symmetric and does not obey the triangle inequality, but it is a measure of loss of information if p is used instead of $\tilde{p}$. The reason we consider $D_{KL}(\tilde{p} \,\|\, p)$, and not $D_{KL}(p \,\|\, \tilde{p})$, is that $\tilde{p}$ is absolutely continuous with respect to p. Since ϑ_i are independent the KL divergence reduces to the sum of KL divergence of the marginal distributions.

In Figure 4a the averaged KL divergence over 100 independent samples is plotted as a function of ρ and is depicted with bars and a scale that corresponds to the left y-axis. For $\rho = 0$ the divergence obtains its maximum while for ρ near 0.2 the divergence is minimized. For ρ larger than 0.3 the divergence slightly increases but remains low. As expected, the mean acceptance rate over the stages of smTMCMC, depicted with filled dots in Figure 4a and scale that corresponds to the right y-axis, drops as ρ increases.

In Figure 4b averaged KL divergence over 100 independent samples is plotted as a function of the number of samples of the sampling algorithm. The KL divergence of TMCMC and smTMCMC with $\rho = 0, 0.2, 0.5, 1$ is estimated. Notice that smTMCMC with $\rho = 0$ shows little improvement as the number of samples increases. This means that the bias introduced by the eigenvalue adaptation

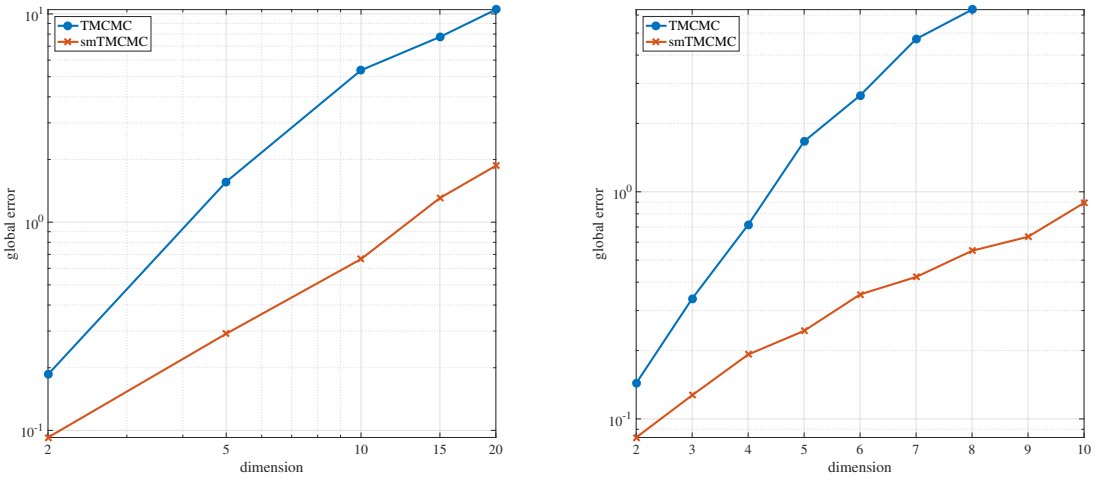

(a) (b)

Figure 5: Averaged error of (a) a unimodal Gaussian distribution (b) a bimodal Gaussian distribution as a function of the dimension.

scheme is not reduced with the size of the sample set. On the other hand, smTMCMC with $\rho = 0.2, 0.5$ and 1 has lower error than the TMCMC, with both slopes being equal to approximately -1. Confidence intervals are not presented because they are smaller than the size of the markers in the plot.

5.2 Multi-Dimensional Gaussian Distribution

In this section we present the scaling of the sampling error for the TMCMC and smTMCMC algorithms in a unimodal and a bimodal Gaussian distribution. In the first test case, the covariance matrix of the target distribution has been randomly generated with the MATLAB function *gallery('randcorr',d)*, where d is the number of dimensions. The error of the sampling algorithm is defined as $E = \frac{1}{2}(e_1 + e_2)$ where $e_1 = ||\bar{\mu} - \mu||_1$, and $e_2 = ||\bar{\Sigma} - \Sigma||_1$ and $\bar{\mu}$ and $\bar{\Sigma}$ are the estimated mean and covariance matrix, respectively. The reported sampling error is averaged over 100 independent simulations using 500 samples. The prior is a d-dimensional uniform distribution with support in $[-10, 10]$ for each dimension. We set the scaling parameters $\varepsilon = 0.04$ for TMCMC and $\varepsilon = 1$ for smTMCMC. Notice that in this case the Hessian matrix is equal to Fisher information and equal to the inverse of the covariance matrix of the target distribution.

In Figure 5a the error for $d = 2, 5, 10, 15, 20$ is presented. The size of the error bars is smaller than that of the markers and thus not included in the plot. Both errors increase with the dimension of the target distribution while the error of smTMCMC is always bellow the error of TMCMC.

Next, we test the performance of the sampling algorithms in a d-dimensional bimodal Gaussian distribution with the two modes centered at $\mu_1 = (-5, \ldots, 5)$ and $\mu_2 = -\mu_1$ and two equal covariance matrices $\Sigma_1 = \Sigma_2$ randomly generated as in the previous example. The scaling parameters are the same as in the previous example. Note that the Fisher information in not explicitly known for Gaussian mixtures, hence the Hessian is the only metric we take into consideration for smTMCMC. To calculate the errors, we first assign each sample to either one of the modes depending on the shortest Euclidean distance. We define the error of the sampling as $E = \frac{1}{4}\sum_{i=1}^{4} e_i$ where $e_i = ||\bar{\mu}_i - \mu_i||_1$ and $e_{2+i} = ||\bar{\Sigma}_i - \Sigma_i||_1$ for $i = 1, 2$ and $\bar{\mu}_i$ and $\bar{\Sigma}_i$ is the estimated sample mean and sample covariance matrix of each one of the two modes.

In Figure 5b the error E averaged over 100 independent simulations is presented. The size of the error bars is comparable to that of the markers and thus not included in the graph. The sample size used in this example is chosen to be $Ns = 5000$. The reason it is increased compared to the unimodal test case is that TMCMC is not able to detect both modes of the bimodal distribution with less samples. Moreover, with the number of samples fixed, TMCMC is able to detect both modes up to $d = 8$ while smTMCMC can go up to 10. Lastly, the difference between the TMCMC and smTMCMC error is increased as a function of dimension.

The computation of the Hessian or Fisher information comes with an additional cost, but during our experiments no remarkable runtime differences have been observed due to the small compu-

tational intensity of the problem. The maximum relative runtime difference has been detected in high dimensions where smTMCMC runs approximately 10% slower that TMCMC.

A feature of the smTMCMC algorithm is that the number of corrections in the covariance matrix (discussed in section 3.1) drop from a very high value (approximately 90%) for early stages to almost zero during the last stages of the algorithm. This is expected, since during the early TMCMC stages the target distribution is close to uniform. This leads to close to zero values for the Hessian or the Fisher information matrices which in turn leads to quite wide proposal distributions that have to be corrected. As TMCMC evolves, the sampled distribution is getting closer to the target distribution and the Hessian or Fisher information has the potential to get improved, at least for locally log-concave functions.

5.3 A pharamacodynamics model

In this section we will use the tumor growth inhibition model proposed in [19] and infer the parameters given data from tumor measurements from different patients. The model consists of a system of ordinary differential equations,

$$\begin{aligned}
\frac{dC}{dt} &= -\varphi_1 C \\
\frac{dP}{dt} &= \varphi_4 P(1 - \frac{P+Q+Q_P}{K}) + \varphi_5 Q_P - \varphi_3 P - \varphi_1\varphi_2 CP \\
\frac{dQ}{dt} &= \varphi_3 P + \varphi_1\varphi_2 CQ \\
\frac{dQ_P}{dt} &= \varphi_1\varphi_2 CQ - \varphi_5 Q_P - \varphi_6 Q_P \\
C(0) = 0,\ P(0) &= \varphi_7,\ Q(0) = \varphi_8,\ Q_P(0) = 0\ .
\end{aligned} \tag{43}$$

We set $Y = (C, P, Q, Q_P)^\top$ and the quantity of interest is $f(t;\varphi) = P(t;\varphi) + Q(t;\varphi) + Q_P(t;\varphi)$.

The tumor is composed of proliferative tissue (P) and quiescent tissue (Q). The treatment, due to the drug administration (C) directly destroys proliferative cells or transforms the quiescent cells to damaged quiescent cells (Q_P). The damaged quiescent cells can either die or repair their DNA and transform into proliferative cells. Here the correspondence of the parameter vector φ to the parameters used in [19] is $\varphi = (KDE, \gamma, k_{PQ}, \lambda_P, k_{Q_PP}, \delta_{QP}, P_0, Q_0)$. We assume a set of 5 measurements $\mathcal{D}_k = \{(t_i, d_i) : i = 1, \ldots, N_d\}$, that correspond to the evolution of the mean tumor diameter (MTD) [19] of 5 patients, and a set of time instances of drug administration $\{\tau_i : i = 1, \ldots, N_\tau\}$ is provided. At $t = \tau_i$ we restart the simulation and set $C(\tau_i) = 1$. In order to reduce the number of parameters to be infered we assume that at time $t = 0$ a measurement, d_1, is provided and thus $\varphi_8 = d_1 - \varphi_7$.

The likelihood function is based on the assumption (2) of normally distributed model error and the augmented parameter vector is defined as $\vartheta = (\varphi^\top, \sigma_n)^\top$. We choose a uniform prior distribution for the parameters, $p(\vartheta) = \prod_{i=1}^{8} \mathcal{U}(\vartheta_i | a_i, b_i)$ with $a = (10^{-2}, 10^{-2}, 10^{-5}, 10^{-5}, 10^{-5}, 10^{-5}, 10^{-5}, 10^{-5})$ and $b = (20, 20, 2.5, 0.3, 0.05, 0.6, 1, 33)$.

Since the prior distribution is uniform, the maximum a posteriori probability (MAP) estimate coincides with the maximum likelihood (ML) estimate,

$$\vartheta_{ML} = \arg\max_{\vartheta} p(\vartheta|\mathcal{D}) = \arg\max_{\vartheta} p(\mathcal{D}|\vartheta)\ . \tag{44}$$

We optimize the log-likelihood for the different data sets, corresponding to different patients, using the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) algorithm [13]. Then, we draw samples from the posterior distribution and compare the maximum log-likelihood in the set of samples with the ML estimate obtained using CMA-ES. We use this test as an indicator to check whether the sampling algorithm is able to sample the high probability area close to the ML point.

In Table 2 we report the estimated ML by the CMA algorithm and compared with the maximum log-likelihood in the sample set of TMCMC and smTMCMC algorithms. The results of pTMCMC are indistinguishable from those of smTMCMC and thus not reported here. The TMCMC algorithm is unable to identify the ML while the smTMCMC is doing better, especially when the size of the sample set is increased. The difficulty the sampling algorithms face when trying to identify the high probability areas, lead us to formulate the hypothesis that an unidentifiable manifold is present in the posterior probability space. Using the the profile log-likelihood (PL) function [18],

$$\mathrm{PL}(\vartheta_i|\mathcal{D}) = \max_{\vartheta_j, i\neq j} p(\vartheta|\mathcal{D})\ . \tag{45}$$

we were able to verify this assumption. The PL function for the ϑ_1 parameter is shown in Figure 9 showing a large area of unidentifiability for approximately $\vartheta_1 > 2$.

No.	CMA-ES	TMCMC $N_s = 1e4$	TMCMC $N_s = 1e5$	smTMCMC $N_s = 1e4$	smTMCMC $N_s = 1e5$
1	-33.06	-38.81	-36.99	-34.89	-33.66
2	-29.91	-39.94	-35.67	-32.11	-30.96
3	-26.24	-28.11	-30.39	-26.79	-26.40
4	-26.69	-29.85	-28.18	-27.39	-26.88
5	-15.68	-35.41	-36.23	-19.42	-17.87

Table 2: The maximum log-likelihood estimate for the pharmacodynamics model (43) obtained using the CMA-ES algorithm, compared with the maximum log-likelihood found in the sample set using the TMCMC and the smTMCMC algorithms.

In Figures 6, 7 and 8 the posterior samples given data for the first patient are presented. In the diagonal the histograms for the 8 parameters are plotted. In the upper triangular part of the plot the pair samples are plotted and in the lower part a smoothed version of the pair marginal histogram is plotted. First, notice that in all the directions TMCMC is being trapped in local maxima of the posterior distribution and is not able to correctly populate the posterior sample space. This is evident from the spikes on the histograms on Figure 7. This unnatural local mass concentration does not appear in the smTMCMC sample set in Figure 8.

We observe that the TMCMC method does not produce consistent samples. The shape of the estimated distribution varies significantly between individual runs of the algorithm. This variability is not treated by increasing the size of the sample set. This is not the case for the smTMCMC algorithm. The distribution presented in Figure 8 does not change between individual runs of the algorithm or if the number of samples is increased.

The location of the ML estimate using CMA-ES and the the smTMCMC algorithm is marked in the diagonal histograms of Figures 7 and 8 with $\times$ and $\circ$, respectively. It can be observed that there is a discrepancy from the CMA-ES estimate in Figure 7, where 10^4 samples were used. In Figure 8 the discrepancy is alleviated by increasing the sample size to 10^5. With this observations, it becomes evident that the high probability area of the posterior distribution are populated.

6 Conclusion

In this paper we have proposed a new sampling algorithm that is a combination of the Transitional MCMC (TMCMC) algorithm with Langevin diffusion transition kernels. Instead of using isotropic diffusion in the TMCMC algorithm, a more advanced proposal scheme is used based on the time discretization of a Langevin diffusion that takes into consideration the shape of the target distribution. This diffusion is taking place in a manifold that is defined either through the Hessian or the Fisher information matrix of the distribution. A typical complication deteriorating the scheme arises from the requirement of positive definite Hessian of Fisher information matrix, which can not be guaranteed for general distributions. Even if the matrices are positive definite some eigenvalues may be close to zero, severely affecting the efficiency of the algorithm. The case of nearly zero eigenvalues is present in systems with unidentifiable parameters.

To alleviate the problems associated with these matrices, we introduced a correction to the scheme, named manifold TMCMC (mTMCMC). We successfully applied the proposed algorithm using a unimodal Gaussian distribution with probability mass concentrated in the boundaries of the prior distribution as well as with a bimodal Gaussian distribution, greatly increasing the sampling quality. Finally, we showcased the ability of the proposed algorithm to sample challenging, multimodal distributions in the presence of unidentifiable manifolds using the posterior distribution in a pharmacodynamics model, and compared its efficiency with respect to existing sampling methods.

The extension of the algorithm for priors other than uniform is also conditionally possible. If the prior distribution is not log-concave then the proposed algorithm is not applicable. In the case of a log-concave prior distribution, one has to redefine the extended boundary technique, presented in Section 3.1, in the absence of a prior with compact support.

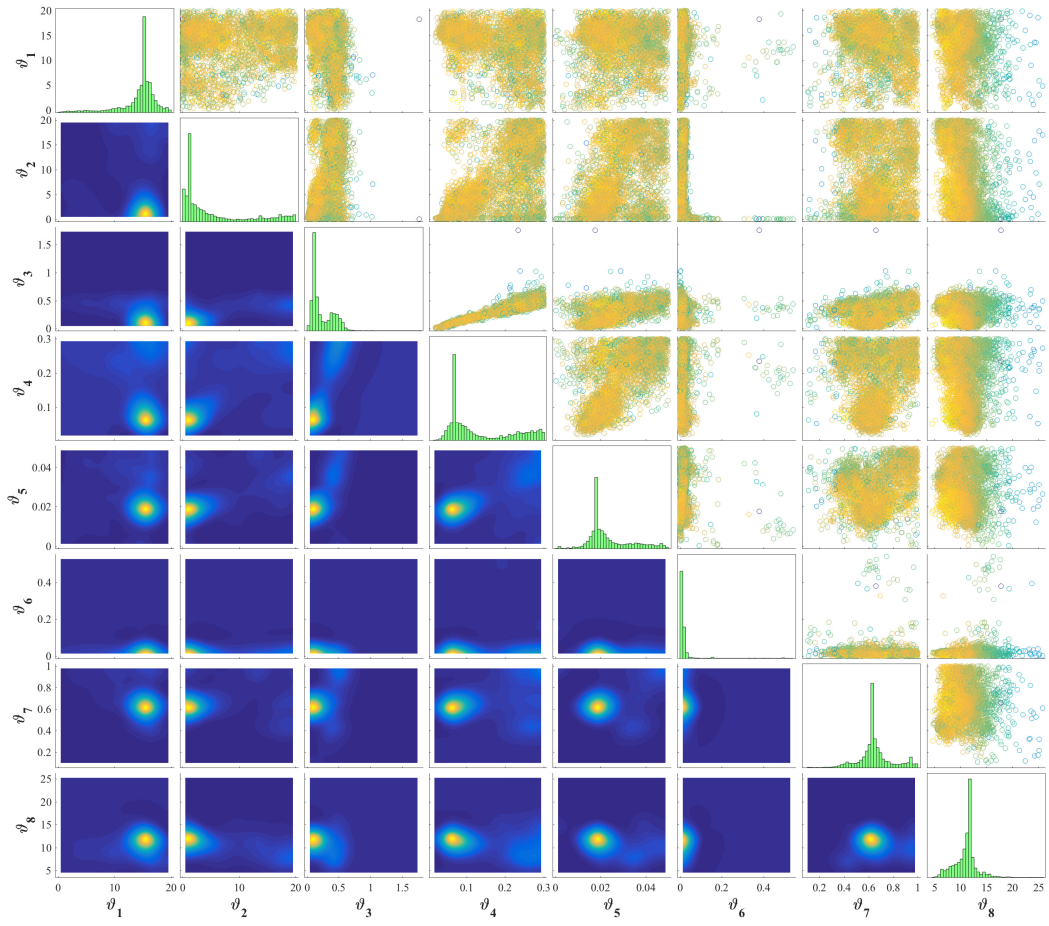

Figure 6: Samples form the posterior distribution of pharmacodynamics model (43) conditioned on data from patient 1 and using the TMCMC algorithm.

7 Acknowledgments

The authors would like to thank Dr Benjamin Ribba and Dr Francois Ducray for permission to use the data in the pharmacodynamics model of Section 5.3.

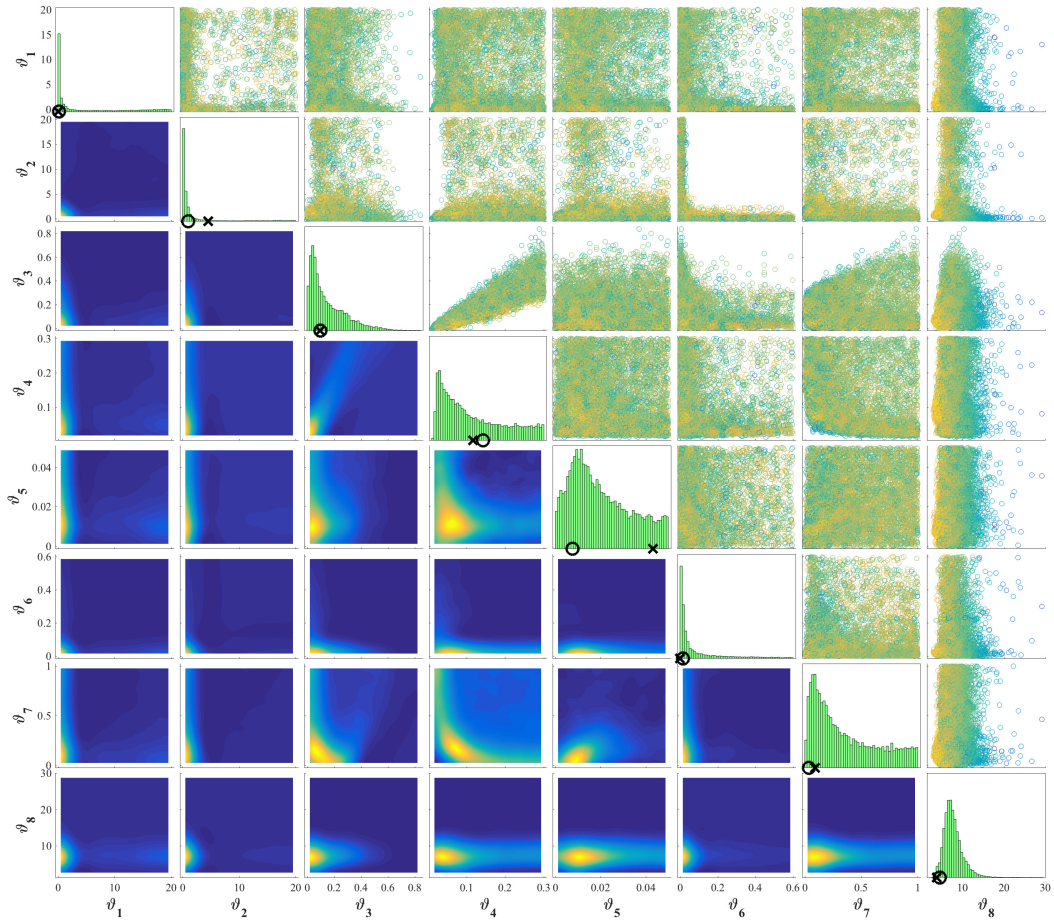

Figure 7: Samples form the posterior distribution of pharmacodynamics model (43) conditioned on data from patient 1 and using the smTMCMC algorithm with 10^4 samples. The $\times$ and $\circ$ correspond to the ML estimates obtained by the CMA-ES and the smTMCMC algorithm, repsectively.

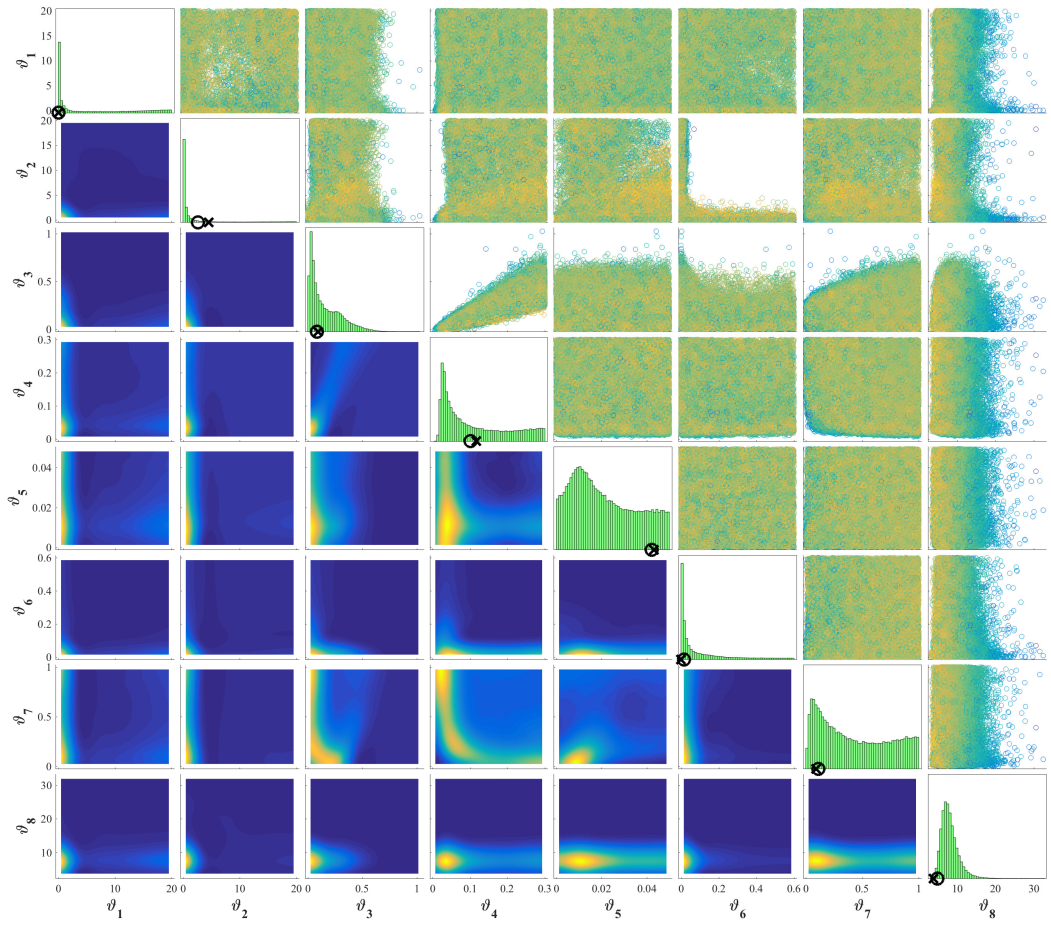

Figure 8: Samples form the posterior distribution of pharmacodynamics model (43) conditioned on data from patient 1 and using the smTMCMC algorithm with 10^5 samples. The $\times$ and $\circ$ correspond to the ML estimates obtained by the CMA-ES and the smTMCMC algorithm, repsectively.

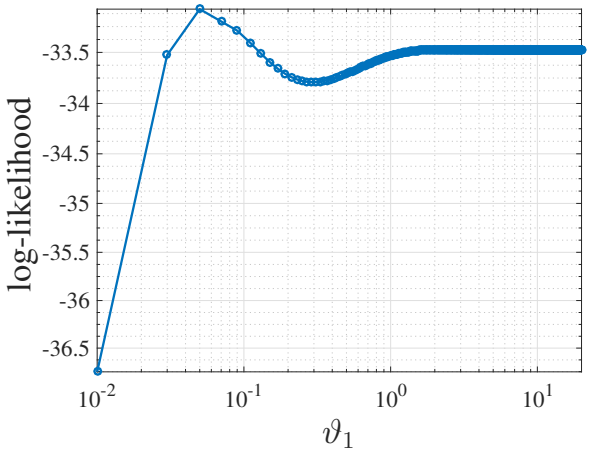

Figure 9: The profile log-likelihood function (45) for the ϑ_1 parameter for the pharmacodynamics model (43) conditioned on data from patient 1.

References

[1] Provided matlab code. `http://cse-lab.ethz.ch/software/`.

[2] K. Aizu. Parameter Differentiation of Quantum-Mechanical Linear Operators. *Journal of Mathematical Physics*, 4(6):762–775, 1963.

[3] J. L. Beck and K. V. Yuen. Model selection using response measurements: Bayesian probabilistic approach. *Journal of Engineering Mechanics*, 130(2):192–203, 2004.

[4] M. Betancourt. *Geometric Science of Information: First International Conference, GSI 2013, Paris, France, August 28-30, 2013. Proceedings*, chapter A General Metric for Riemannian Manifold Hamiltonian Monte Carlo, pages 327–334. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[5] W. Betz, I. Papaioannou, and D. Straub. Transitional Markov Chain Monte Carlo : Observations and Improvements. *Journal of Engineering Mechanics*, pages 1–10, 2016.

[6] T. Bui-Thanh and O. Ghattas. A scaled stochastic Newton algorithm for Markov Chain Monte Carlo simulations. *SIAM Journal on Uncertainty Quantification*, pages 1–25, 2012.

[7] B. Calderhead and M. Girolami. Estimating Bayes factors via thermodynamic integration and population MCMC. *Computational Statistics and Data Analysis*, 53(12):4028–4045, 2009.

[8] J. Ching and Y.-C. Chen. Transitional Markov Chain Monte Carlo Method for Bayesian Model Updating, Model Class Selection, and Model Averaging. *Journal of Engineering Mechanics*, 133(7):816–832, 2007.

[9] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987.

[10] M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 73(2):123–214, 2011.

[11] P. J. Green, K. Łatuszyński, M. Pereyra, and C. P. Robert. Bayesian computation: a summary of the current state, and samples backwards and forwards. *Statistics and Computing*, 25(4):835–862, 2015.

[12] P. Hadjidoukas, P. Angelikopoulos, C. Papadimitriou, and P. Koumoutsakos. Π4U: A high performance computing framework for Bayesian uncertainty quantification of complex models. *Journal of Computational Physics*, 284:1–21, 2015.

[13] N. Hansen, S. D. Müller, and P. Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary computation*, 11(1):1–18, 2003.

[14] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

[15] T. S. Kleppe. Adaptive step size selection for Hessian-based manifold Langevin samplers. *Scandinavian Journal of Statistics*, (1):1–24, 2015.

[16] N. Lartillot and H. Philippe. Computing Bayes Factors Using Thermodynamic Integration. *Systematic Biology*, 55(2):195–207, 2006.

[17] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. A Stochastic Newton MCMC Method for Large-Scale Statistical Inverse Problems with Application to Seismic Inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012.

[18] A. Raue, C. Kreutz, F. J. Theis, and J. Timmer. Joining forces of Bayesian and frequentist methodology: a study for inference in the presence of non-identifiability. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984), 2012.

[19] B. Ribba, G. Kaloshi, M. Peyre, D. Ricard, V. Calvez, M. Tod, B. Čajavec-Bernard, A. Idbaih, D. Psimaras, L. Dainese, J. Pallud, S. Cartalat-Carel, J. Y. Delattre, J. Honnorat, E. Grenier, and F. Ducray. A tumor growth inhibition model for low-grade glioma treated with chemotherapy or radiotherapy. *Clinical Cancer Research*, 18(18):5071–5080, 2012.

[20] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.

[21] G. Roberts and O. Stramer. Langevin diffusions and Metropolis-Hastings algorithms. *Methodology And Computing In Applied Probability*, 4(4):337–357, 2002.

[22] G. O. Roberts and J. S. Rosenthal. Optimal scaling of discrete approximations to langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268, 1998.

[23] G. O. Roberts and J. S. Rosenthal. Optimal Scaling for Various Metropolis-Hastings Algorithms. *Statistical Science*, 16(4):351–367, 2001.

[24] M. Slotani. Tolerance regions for a multivariate normal population. *Annals of the Institute of Statistical Mathematics*, 16(1):135–153, 1964.

[25] S. Wu, P. Angelikopoulos, C. Papadimitriou, and P. Koumoutsakos. Bayesian Annealed Sequential Importance Sampling (BASIS): an unbiased version of Transitional Markov Chain Monte Carlo. 0(0):1–19, 2015.

[26] T. Xifara, C. Sherlock, S. Livingstone, S. Byrne, and M. Girolami. Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Statistics and Probability Letters*, 91(1):14–19, 2014.

A Derivative of the transformed matrix

Proposition 1. *Let $A \in \mathbb{R}^{n\times n}$ be a matrix that depends on a parameter ϑ and let $A = Q\Lambda Q^\top$ its eigendecomposition. Let $f(A) = Qf(\Lambda)Q^\top$ be a transformation of the matrix A. Then, it holds that*

$$\frac{\partial f(A)}{\partial \vartheta} = Q\Big(J \circ \big(Q^\top \frac{\partial A}{\partial \vartheta} Q\big)\Big)Q^\top \,, \tag{46}$$

where $\circ$ denotes the Hadamard product and J is given by,

$$J_{i,j} = \begin{cases} \frac{f(\lambda_i)-f(\lambda_j)}{\lambda_i-\lambda_j}, & i \neq j \\ \frac{\partial f(\lambda_i)}{\partial \lambda_i}, & i = j \,. \end{cases} \tag{47}$$

Proof. For the proof see Section 2 in [2]. □